UNCLASSIFIED

Defense Technical Information Center Compilation Part Notice

ADP014132

TITLE: The Assessment of Engine Usage Data

DISTRIBUTION: Approved for public release, distribution unlimited
Availability: Hard copy only.

This paper is part of the following report:

TITLE: Aging Mechanisms and Control. Symposium Part A - Developments in Computational Aero- and Hydro-Acoustics. Symposium Part B - Monitoring and Management of Gas Turbine Fleets for Extended Life and Reduced Costs [Les mecanismes vieillissants et le controle] [Symposium Partie A - Developpements dans le domaine de l'aeroacoustique et I'hydroacoustique numeriques] [Symposium Partie B ...

To order the complete compilation report, use: ADA415749

The component part is provided here to allow users access to individually authored sections of proceedings, annals, symposia, etc. However, the component should be considered within the context of the overall compilation report and not as a stand-alone technical report.

The following component part numbers comprise the compilation report:
ADP014092 thru ADP014141

UNCLASSIFIED

The Assessment of Engine Usage Data

D.P. Shepherd
Structures and Materials Centre, QinetiQ
Farnborough, Hants, GU14, 0LX, UK
Phone +44-1252-392000
Fax +44-1252-397298
Email: dpshepherd@qinetiq.com

1 INTRODUCTION

Within aero gas turbine engines, fracture critical parts are defines as those whose failure could hazard the entire aircraft. Because of the potentially catastrophic impact of such a failure, a fundamental airworthiness requirement on all aero gas turbine engines is that the operational life of these parts is managed to ensure this does not happen. Consequently, considerable effort is devoted to deriving methods for estimating the usable life of such components, to ensure that the probability of in-service failure is acceptable low. Such methods must accurately account for the basic material properties, the impact of the various manufacturing processes employed, and the operational loads the components will be subjected to during service usage. In general terms, the lives are derived on the basis of laboratory specimen test and/or full scale component rig tests, finite element stress analyses, together with appropriate statistical models of fatigue behaviour. A vigorous and ongoing research activity is devoted to developing and refining these methods, and to ensure that they are kept up to date with changes in the basic engine technology.

From the available information, these estimation procedures will produce component lives quoted in terms of some reference stress cycle. Once this has been done, however, there remains the problem of relating these reference cycle lives to operationally measured parameters. For civil applications, there is only minor throttle movement once the aircraft has reached cruise conditions, and the engines see a standard, relatively simple operational cycle. For this reason, reference cycles can fairly accurately be related to flights on a one-to-one basis. However, for military operation, the problem is considerably more complex. In many military sorties, the output from the engines requires constant adjustment, which inflicts considerable amounts of minor cycle damage on the engine components. Moreover, the level of this minor cycle damage will itself vary considerably from mission to mission, even for the same type of sortie. Thus, it is essential to account for these fluctuations in usage to ensure adequate life consumption tracking.

The issue of relating complex mission profiles to the consumption of reference cycles involves two separate technical problems. The first of these is the physical monitoring of how the engine is being used in practice, and the recording of this information. This is achieved by fitting the engines with instrumentation that measures a number of important parameters, such as spool speeds, gas temperatures and so on, which are recorded on data logging equipment. The second problem involves the development of engine usage algorithms, which interpret the recorded data and provide an assessment of the resulting usage. These algorithms translate the recorded parameters into mechanical and thermal loads, which are then converted to reference cycles through damage accumulation models. Using these techniques, it is possible to provide an equivalent reference cycle count for any mission, no matter what the complexity of the manoeuvres undertaken.

However, the instrumentation of the engines, together with the collection and management of the data generated, represent considerable costs to the operator. For this reason, it has traditionally been the case (at least in the UK), that only a sample of engines from the fleet have been monitored in this way. Whilst this approach clearly represents an immediate cost benefit, it introduces a third element in the life tracking process; namely, that of relating the data on life consumption from the monitored aircraft to that of the rest of the fleet. To do this, it is necessary to exploit the fact that a number of parameters relating to the operation of the aircraft, such as sortie type and duration, are routinely recorded for all missions part of the fleet management process. By developing a relationship for the monitored engines between cycles

Paper presented at the RTO AVT Symposium on "Ageing Mechanisms and Control: Part B – Monitoring and Management of Gas Turbine Fleets for Extended Life and Reduced Costs", held in Manchester, UK, 8-11 October 2001, and published in RTO-MP-079(I).

recorded and flight parameters, cycle count predictions can be estimated for all unmonitored situations. However, this introduces a further difficulty, due to the fact that the damage recorded from individual flights can vary significantly, even under nominally identical conditions. Thus, the relationship between the cycles consumed and the flight parameters is probabilistic in nature, and a statistical model must be developed to represent this relationship. Moreover, the manner in which this model is related to the rest of the fleet must also, therefore, be statistical in nature.

In this paper, a statistical model for the analysis of engine usage data generated from a sample of monitored aircraft is developed, and its application to a set of cycle counts is described. In section 2, the UK experience in military aircraft monitoring is described, and the calculation methods used in the application of these data are discussed. Next, a set of usage data, taken from a current UK military engine programme, is introduced, and the results from an initial analysis are illustrated. An appropriate statistical model describing the relationship between consumed cycles and relevant operational parameters is then developed, and fitted to the data. Current and future practice with regard to the prediction of life consumption for unmonitored aircraft is considered in the light of this. Finally, areas where further work is required are highlighted.

2 HISTORICAL BACKGROUND AND CURRENT PRACTICE IN THE CALCULATION OF EXCHANGE RATES

For the situation in which only a subset of engines from a particular fleet are instrumented to record accumulated damage, the fundamental problem in tracking component life is to provide an estimate of cycles consumed for those engines which are not monitored. This requires that a prediction of consumed cycles be derived solely on the basis of the available information which describes the unmonitored missions. The purpose of developing a model from the monitored missions, therefore, which describes the relationship between these measured parameters (referred to below as explanatory variables) and the actual cycles consumed, is to provide this predictive capability. However, to provide reliable and accurate predictions, it is necessary to account for any relationships which may exist between the measured values and the descriptive parameters. For example, suppose that a particular mission flown from one operational base results in higher average levels of accumulated damage than the nominally identical mission flown from a second base. If the two different averages can be successfully identified, then these can be used separately to provide predictions of the behaviour for future missions flown from these two bases. However, if a simple average of the total set of missions from both bases were used to provide the predictions, usage at the first base would be systematically underestimated, and that at the second overestimated. Such systematic discrepancies in cycle count predictions could lead to significant errors over the lifetime of a component, and in the worst case could result in a serious underestimation of life consumed. From an airworthiness point of view, therefore, it is important to identify any underlying relationships between the explanatory variables and the measured cycle counts, so that these can properly be taken into account. The problems to be overcome in fulfilling this task are, firstly, that such trends can be difficult to identify where several influencing variables are acting simultaneously, and secondly, the statistical scatter inherent in the data will tend to mask further these underlying trends.

In the UK, it has traditionally been considered that the dominant explanatory variable influencing the accumulation of cycles is mission duration, and that usage is assumed to be a direct multiple of the flight time. This dependence has been expressed in terms of 'cyclic exchange rates' (also referred to as 'β' factors), which are calculated for any mission simply by dividing the cycles consumed by the duration. The required relationship is defined by identifying a representative exchange rate, based on the individual exchange rates calculated from each monitored mission. In very early exchange rate calculation exercises, the representative β factor was often calculated simply as a gross average of measured exchange rates. The influence of other systematic variables on the exchange rate was accounted for by adding various 'safety factors' to this average, with little or no attempt to directly quantify the underlying relationships. Moreover, the means by which these additional safety factors were calculated varied between the different engine projects (ref 1). However, as time progressed, either role or sortie pattern code were introduced as additional explanatory variables. In these cases, the data was broken down according to the variable being used, and individual averages were established for each case. A fleetwide exchange rate was then established by estimating the (average) time spent in each role/SPC, and using this to calculate a weighted average of the individual β's (ref 2).

In practice, these methods were found to give acceptable results, as long as the type of flying undertaken by any particular engine application was reasonably simple and stable. However, as operational requirements became more complex, particularly with the introduction of multi-role aircraft, some doubt as to the adequacy of this approach was raised. As a result, in the mid 1970's, the RAF enlisted the aid of the MoD's statistical branch, Stats (S) 3, to provide a more robust solution to the problem. Over the next several years, a detailed statistical model of the then available data was developed, and a series of computer programmes were written to automate the exchange rate calculation process (ref 3). This approach, which became known as the 'Baillie method' after the author, made a direct attempt to assess which operational variables had a significant impact on the overall consumption. The statistical model proposed by Baillie (in the single engined case), was expressed by the equation

$$y_{jkr} = A + E_j S_k t_{jkr} \varepsilon_{jkr}.$$

In this equation, A is a factor representing the cycle consumed during take off and landing within each flight, E_j is a multiplicative factor representing the effect of the jth engine, S_k accounts for the effect of the kth Sortie Pattern Code/role combination, t_{jkr} is the duration of the rth sortie flown within the particular engine/SPC combination, and ε_{jkr} is a lognormally distributed random variable, which accounts for the random variation. From this equation, it can clearly be seen that usage is not assumed to be a simple multiple f the flight time. Rather, the two are related according to a straight line which is offset by a constant factor (A), and hence does not pass through the origin. Moreover, engine and SPC are explicitly included as explanatory variables in the model. Estimates for the required constants were obtained from the data through a maximum likelihood approach. A safe exchange rate for the fleet was then derived from this analysis, incorporating an additional factor to account for the worst likely engine in the fleet. This was obtained from the individual engine effects, E_j, by fitting a lognormal distribution to them and taking an upper percentage point appropriate to the number of engines in the fleet. This method was used to derive cycle count predictions for the Spey engines powering the RAF's fleet of Phantom aircraft.

However, the application of this method on one particular project led to further concerns regarding the lack of commonality between engines, and attempts were made to derive a common method suitable for all applications (ref 2). Despite (or perhaps because of) its statistical rigour, the Baillie method was rejected as being too complex, and a simpler alternative was proposed (ref 5). This was based on fitting a linear regression line to plots of cumulative cycles against cumulative flying hours, and computing a suitable upper confidence bound. Unfortunately, this too failed to produce the desired commonality, and different methods continued to be used on a number of programmes. In the early 1990's, the Weibull distribution was adopted within some engine projects as offering a superior description of exchange rate distributions to the lognormal. However, the analysis still consisted of breaking the data down into individual SPC's, fitting distributions and deducing representative rates (by going to an upper percentile to account for systematic variability), and calculating a weighted average. Despite the enormous advances that have been made with regard to damage counting algorithms, and the sophistication of modern data logging systems, it remains the case that the exchange rate calculation procedure itself is based on a very simplified representation.

3 DATASET AND INITIAL DATA ANALYSIS

To investigate the assumptions which form the basis of the various exchange rate calculation procedures discussed above, a set of cycle counts for a particular engine have been examined. To do this, the data have been subjected to an initial visual inspection, to try and assess whether any clear relationships can be identified. A statistical model has then been constructed and fitted to the data, so that these assumptions can be examined in more detail. This leads to a number of conclusions regarding current exchange rate calculation practice, and suggestions for areas in which further analysis may provide significant improvements.

The data used in the current analysis are taken from a current UK military engine programme. The total dataset consists of nearly 3000 individual engine recordings, taken over a period of 14 years. The aircraft is twin engined, and data are recorded[1] for both engines separately, so the actual number of flights from which the data are taken is considerably less than this. (It is greater than half, though, since there are a

[1] The sampling rate used is 8 Hertz

significant number of cases where data from one or other of the engines is missing). For each flight, 9 parameters relating to the flight are manually recorded, in addition to the cycles consumed. These are aircraft tail number, sortie number, engine number, mark of engine, whether equiaxed or single crystal turbine blades were fitted, the base from which the mission was flown, the date, duration of the flight and sortie pattern code. The cyclic damage consumption for a single component within the engine is considered, though it is strongly anticipated that the structure of the data will be very much the same for all other components.

As with any statistical model fitting exercise, the first step is to have an initial look at the data using knowledge of the situation and simple graphing techniques. The first point to establish is which variables are likely to be the important ones in developing an adequate statistical model. Common sense plus a knowledge of previous exchange rate models clearly indicate that SPC, base and duration are likely to be the primary correlating variables. Whilst tail number, sortie number and engine number may be of use in exploring detailed structure within the data, they clearly play little or no role in determining cyclic consumption in themselves. Similarly, whilst it may be of interest for some purposes to consider whether there are any trends in the data with respect to time, this is not considered within the present study. Within the sample, all engines are of the same mark, so this is clearly irrelevant from the point of view of statistical modelling. For the type of turbine blade, it is known that this does have an effect, since an allowance for the extra load induced by the heavier single crystal blades is built into the damage calculation algorithm. However, the effect is small (5-10%) and predictable, so it is considered to be very much a second order effect.

A cursory inspection of data immediately reveals that it is very uneven in terms of the number of recorded flights within any particular category. For example, the data contains recorded flights from 40 different bases, of which 29 have less than 10 recorded flights. Similarly, there are 34 different sortie pattern codes recorded in the data, of which 23 contain less than 20 flights. Since it will be virtually impossible to deduce anything of value regarding the main effects or there interactions with so little data, these very weakly populated factors have been removed from the analysis. This leaves 11 sortie pattern codes and 11 bases to consider. Even within this reduced dataset, however, the number of flights in each base/SPC combinations is highly variable, and there are a number which do not contain any data at all (see Figure 6). Clearly, this will have a influence on the information which can be derived regarding engine usage in these particular situations, and in practice it makes the statistical modelling process considerably more complex (see below).

Using this reduced dataset, one of the first points to consider is whether the distribution of cycles per hour (or cycles) is best described using a Weibull or lognormal model. The easiest and most reliable way to do this is to look at the quantile-quantile[2] plots for different SPC's and bases. Figures 1-4 clearly show that a lognormal distribution provides a superior description of the data. Indeed, points that appear as outliers in the Weibull plot of Figure 3 are seen to fall into the body of the distribution in the lognormal plot of Figure 4. Consequently, in the following analysis, the random element in the model is assumed to follow a lognormal distribution, (as was used in the Baillie approach). This means that, by taking a log transformation of the data, the normal distribution becomes the basis for the model. This has interesting consequences from the statistical point of view, since it means that the classical least squares regression and analysis of variance techniques can be used. The consequences of this are discussed in the next section.

A further point that can usefully be addressed at this stage is the regression of the cyclic damage counts against the mission duration. Figure 5 show multiple plots of the log cycles per hour and log cycles, broken down by both SPC and base. It can be seen that the regression against log cycles is somewhat variable in nature, in some situations appearing to be virtually flat[3]. This suggests that the current practice of calculating cycles consumed as a multiple of flight duration is somewhat unrepresentative of actual experience, and could lead to a significant underestimation of the consumed damage for short duration flights. Moreover, it suggests that the approach adopted by Baillie is much more representative in this respect also, since it allows for regression lines which do not pass through the origin. Notice also that the very flat nature of the regression response in certain situations suggests mission duration has

[2] A quantile-quantile plot (or Q-Q plot), is a plot of the ordered values against the corresponding quantiles of the respective distribution

[3] This observation is also confirmed by other studies of cycle count data (see ref 8)

virtually no influence on accumulated damage in these cases. However, the plots still suggest that at least some of the regression coefficients will be non-zero, indicating that it will still be necessary to include mission duration as an explanatory variable.

4 MODEL FITTING

Once a preliminary visual examination of the data has been conducted, the next step in any statistical analysis is to examine the basic observations already established, and to explore any further relationships which may not be so immediately obvious. This is achieved by fitting statistical models to the data, which allows for multiple causal relationships to be explored, and which ensures that the random variation present in the data to be properly accounted for. The process involves constructing statistical models containing multiple variables, and assessing their validity by examining the extent of the random variation still present once the fit has been achieved. More specifically, in the presence of random variability in the raw data, the influence of any particular variable will be revealed through the extent to which this variability is partitioned. For example, suppose that missions flown from a particular base consumed significantly more damage, on average, than the equivalent missions at another base. Whilst the data from the two bases may well overlap, the scatter associated with each individually will be significantly less than that that produced by both combined. Thus, the greater the extent to which the variability is reduced in fitting a model containing that variable, the more significant the variable is in 'explaining' the data. The aim of the statistical analysis is to reduce the extent of the scatter as far as possible, by accurately identifying the systematic relationships within the data.

It is in this sense that the ability to assume a normal description for the underlying variability is significant. It was mentioned in section 3 that this assumption allows the classical methods associated with least squares analysis to be employed, and in particular techniques generally referred to as the Analysis of Variance. This is significant, because this is the branch of statistics which was developed specifically to investigate problems involving the partitioning of variability (as the name implies). It is also significant, because potentially it represents a distinct advance over the methods used in the Baillie analysis. One reason for this it that it allows for the construction of specific hypothesis tests, which can be used to assess the relative significance of the different variables. The reason Baillie himself did not take advantage of these techniques is probably due, at least in part, to the limitations in computer power available at the time. The basic solution techniques, when applied to datasets of the size considered here, requires the inversion of very large matrices which would have swamped most computer systems available in the 1970's. However, given that todays computers can perform such tasks as a matter of routine, the very powerful techniques developed to examine this type of problem can now be applied to the analysis of cycle count data.

In spite of the advances in computer technology, however, there are still significant difficulties involved in applying this type of analysis. The most important of these is the fact that the data points are distributed very unevenly amongst the cells of the SPC/base categorisation, as observed in section 2. This is unfortunate, because it severely affects the type of hypothesis tests which can be conducted, and under the regular form of analysis it renders them largely meaningless. More sophisticated methods for dealing specifically with this type of data do exist (see ref 8), but require a significant effort to employ because they are not generally available within commercial statistical analysis packages[4]. Since this was not possible within the scope of the present study, such test have not been used. A further consequence of the unbalanced nature of the data, and particularly the fact that some cells are empty, is that there are certain quantities within the model which cannot be estimated. This should not be surprising, since it is clearly impossible to derive meaningful conclusions concerning a physical phenomenon with no available data. However, is does illustrate the importance of increasing the available sample size to ensure that all operational situations are being appropriately accounted for.

The model fitted to the data in the current analysis is of the form

$$\log C_{ijk} = \alpha + S_i + B_j + S_i * B_j + Dx + \varepsilon_{ijk}$$

---

[4] This may have been another reason why Baillie did not use these techniques, since the methods required to analyse this type of data were in the process of being developed at the time.

where S is SPC, B is base and D denotes log mission duration. Moreover, the same model has been fitted repeatedly, but with some of the terms removed, to examine the influence of each term separately. An example of the output from the fitting programme is shown in figure 7, displaying the calculated constants. From the statistical modelling point of view, the important number in figure 7 is the sum of squares due to the residuals, which is given as 141.7. This is important, because it represents the amount of random variability left after the model has been fitted. By examining this figure as the different terms from the model are removed, the effectiveness of each term in reducing this variability can be assessed. This process reveals that this figure does not vary by large amounts, as shown in the following table[5].

Terms included in model	Residual sum of squares
Base, Hours, SPC, Base*SPC	141.7
Base, Hours, SPC	153.5
SPC, Hours	160.0
Base, Hours	175.9
Hours	187.5
None	192.9

The conclusion to be drawn from this is that effect the explanatory variables employed in the model is largely swamped by the residual variability, which confirms the visual analysis given in Figure 5. The remaining plots in Figure 8 to Figure 10 show the residuals (that is, the difference between the model and the observed values at any point) plotted in various ways, and are diagnostic in nature. They show that, essentially, none of these assumptions used in constructing the model are violated. In particular, the histogram in Figure 8 shows hat the assumption of underlying normality is well justified. The conclusion from these plots is that the model fits extremely well, and provides a more than adequate representation of the data.

5 DISCUSSION

Before moving on to consider the implications of this analysis for cycle count prediction methods, some features of the analysis warrant further discussion. For example, the fact that the mission duration shows little or no effect on the cyclic damage accumulation in under certain sortie and base combinations may at first appear surprising. However, for many types of flying undertaken by the RAF, there are severe limits imposed on how this can be carried out. For example, there are relatively few places in Britain where low level flying activities can be undertaken. Consequently, aircraft from different bases will need to use the same airspace for this activity, and the transit times to and from these locations to the bases themselves will vary significantly. Consequently, whilst the mission times will be very different depending on which base they are flown from, it is reasonable to suppose that the actual level of damage will be more or less the same. Thus, when one considers the constraints under which many of the missions are flown, it starts to appear that this is to be expected after all.

A natural consideration which immediately arises from the relative lack of success in partitioning the variability, is whether better correlating factors can be found. In fact, there are a number of other operational parameters routinely recorded against flights of these engines, which may perform this function. Such factors include number of landings, number of in-flight refuellings and role. Whilst these may allow for a better description of the variability, it seems likely that they will only partially explain the observed scatter. Previous studies have indicated strongly that there is a large variability in the way that different pilots fly combat aircraft, and without explicitly including this effect into the analysis, it is unlikely that the observed scatter in consumed cycles can be reduced much beyond that achieved in the

[5] It is important to remember that these figures represent the sums of squares of residuals, so the percentage change is much smaller than represented here.

current analysis. Consequently, it appears likely that this feature of the data will remain a dominant influence in determining overall exchange rates.

Next, there is a further feature of the data presented here which, if modelled correctly, could provide very significant benefits in terms of both exchange rate calculation and general fleet management. Figure 11 shows a plot of a series of cycles counts, taken over consecutive sorties flown by a single aircraft between engine removals. It shows that the right hand engine consistently accumulates greater amounts of damage than the left, irrespective of the severity of the sortie. The reason for this discrepancy is the variation of build between the two engines, which leads to a differential between the spool speeds and temperatures required to provide the same amount of thrust. It is well known that such differences exist in service engines, and it is not itself suprising to observe such differences being recorded in practice. However, there would be many advantages to be gained if these variations could be quantified in such a way that they could be compared between tail numbers. For example, this would allow for a general distribution of engine variability to be derived, which would quantify the overall spread in build standard. Moreover, by the process of extracting this information would effectively remove the contribution made by build quality to the overall variability. Finally, provided enough raw information was available, it could be used as a diagnostics tool to assess trends in build quality over time, and ensure that common standards were being maintained. The difficulty with Figure 11 as it stands is that it only allows for the variation in these two engines to be assessed relative to each other. In order to provide a comparison between engines in different aircraft, they would need to be assessed relative to the mean damage rate for that particular combination of explanatory variables, taking into account the residual scatter. This would require a more sophisticated model, which would included the effect of engine in an appropriate way. Notice also, that this is precisely the purpose of the parameter E_j in the Baillie model. This emphasises further the comprehensive nature of the Baillie analysis, given the constraints that were operating at the time. Finally, it is important to reflect on the implications of the analysis on current usage prediction and life consumption monitoring practice. Whilst the introduction (or, more accurately, re-introduction) of the lognormal distribution for describing the underlying variability in the data has a very significant impact in terms of the models used, in itself it is unlikely to have a large effect on the numerical estimates obtained. The relative flatness of the regression lines when plotting log cycles against mission, on the other hand, suggest that the current practice of multiplying the number of cycles consumed per hour by the mission duration could introduce significant errors into the life management process. The importance of this error, however, can only be assessed against the extent to which the engines experience a random sample from the population of consumed cycles. Indeed, if engines really do sample at random from the distribution, then it is unlikely to have a significant impact overall, since the underestimation at low mission durations will be compensated for by a parallel overestimation for long missions. But if this really is the case, then the current practice of taking an upper percentage point on the distribution as defining a safe exchange rate for the fleet would appear to be overconservative. This is because if the sampling is actually random, the large number of sorties flown by an engine during its service lifetime will mean that its cyclic life consumption per mission, averaged over its life, will be very close to the mean. Consequently, the whole question of the most appropriate means of estimating fleetwide exchange rates depends on the investigation of this problem. Thus, the work following on from the current analysis will be to investigate this problem, so that the most accurate and cost effective means of exchange rate calculation can be determined.

6 CONCLUSIONS

The analysis of low cycle fatigue cyclic consumption data for military aircraft has been reviewed, for the situation in which only a sample of engines are monitored. In particular, the analysis developed by D H Baillie has been described, and contrasted with past and present methodology. Moreover, the problems to be addressed in this type of analysis have been identified, and some assumptions underlying current practice are described. To investigate the validity or otherwise of these assumptions, a set of cycle count data from a modern UK military aircraft has been analysed. Initial investigation of the data reveals that previous assumptions concerning the probability model describing the scatter in the data are not optimal, and a lognormal distribution appears preferable to the commonly used Weibull. This has a distinct advantage form the statistical point of view, since it means that the classical methods relating to the analysis of linear models can be applied directly to the problem. Consequently, a regression model has been derived and fitted to the data, which includes two factors in addition to mission duration. The analysis, however, is complicated by the unstructured nature of the data, and particularly the fact that the observations are very unevenly distributed through the classification scheme. The results indicate that whilst the model fits the data perfectly adequately, the explanatory variables employed in the analysis only provide a relatively weak description of the data. In the course of the analysis, however, three areas of further work are identified, which may allow substantially stronger results to be derived. These are,

1) The development of an analysis capability which can fully account for the unbalanced nature of the data, resulting in the full power of the available techniques being applied.
2) The introduction of additional explanatory variables into the model, which could potentially provide a much improved description of the data.
3) The development of techniques which could quantify the effect of variation in build quality on the consumption rate, which could provide substantial benefits to both cyclic consumption prediction and fleet management.

Finally, the problem of applying this type of analysis to the problem of fleetwide usage calculation has been addressed, and the need for a comprehensive analysis of the operational data for the fleet is emphasised.

The author would like to thank the support of the MoD in providing the funds for the work to be carried out, and to the RAF for providing the data.

7 REFERENCES

1) J A Lang, 'Towards an Agreed Method of Estimating a Safe Cyclic Exchange Rate From EUMS/EDR data', D/D Eng D(PE)/19/61/2, 28 Sept 1982.
2) 'Minutes of a Meeting Held to Discuss the derivation of Military Engine Low Cycle Fatigue Exchange Rates', D/D Eng T(PE)/4/33/9, 6 Dec 1982
3) D H Baillie and G J Knightley, 'Rederivation of Aero-Engine Low Cycle Fatigue Consumption Rates', Defence Statistics report, Nov 1976
4) D H Baillie, 'A Statistical Model of Low Cycle fatigue Consumption in Single Engined Aircraft', Defence Statistics report, Oct 1980
5) R H Jenkins, 'The Derivation of Cyclic Exchange Rates Calculated From Flight Recorded Data', Rols-Royce report BSR 4789, 9 Dec 1983
6) T M Edmunds, 'Derivation of Exchange Rates From Flight Recorded Data', Rolls-Royce report BSN/D/1039, 6 Jun 1984
7) H Pfoertner and J Broede, 'Statistical Correlations for Analysis and Prediction of Turbine Engine Component Life Usage', 37th AIAA/ASME/SAE/ASEE Joint Propulsion Conference, Salt Lake City, Utah, July 2001
8) G A Milliken & D E Johnson, 'Analysis of Messy Data: Volume 1, Designed Experiments', Van Nostrand Reinhold, 1984

8 FIGURES

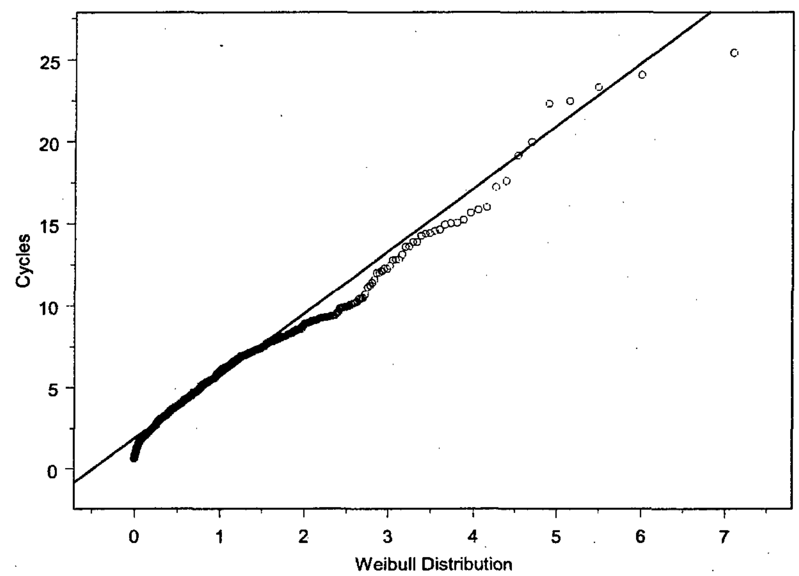

Figure 1 Fit of Weibull distribution to SPC data

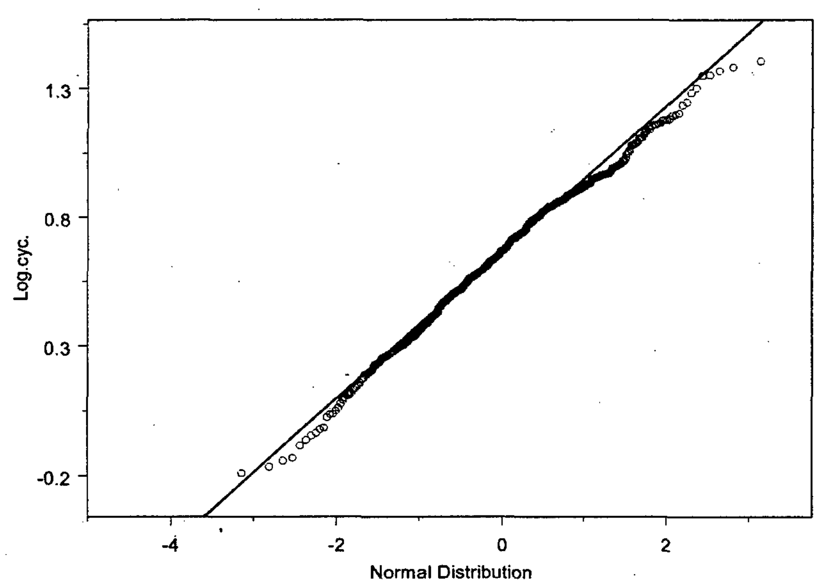

Figure 2 Fit of lognormal distribution to SPC data

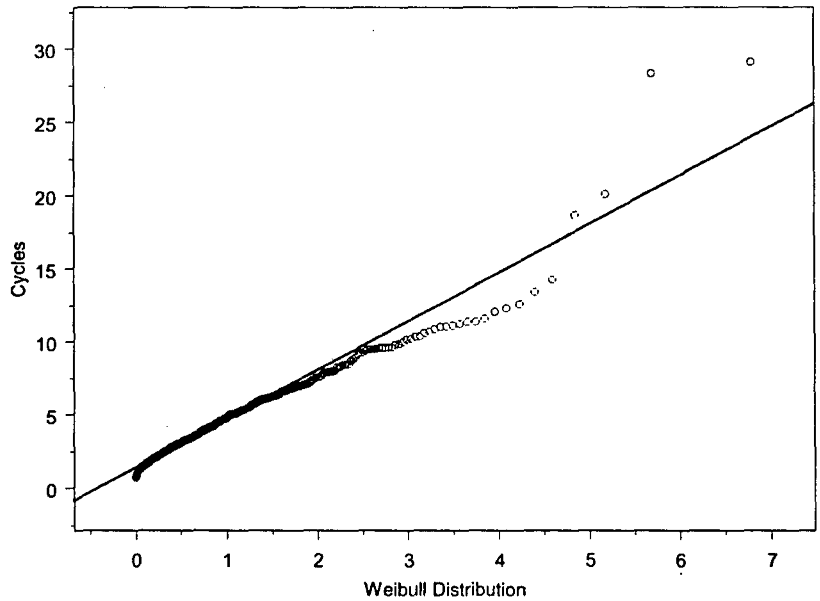

Figure 3 Fit of Weibull distribution to SPC data

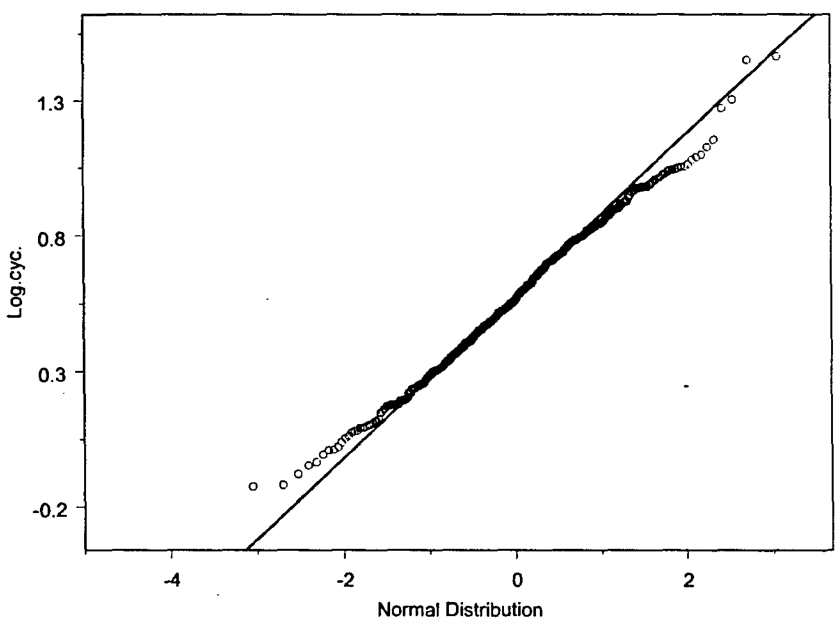

Figure 4 Fit of lognormal distribution to SPC data

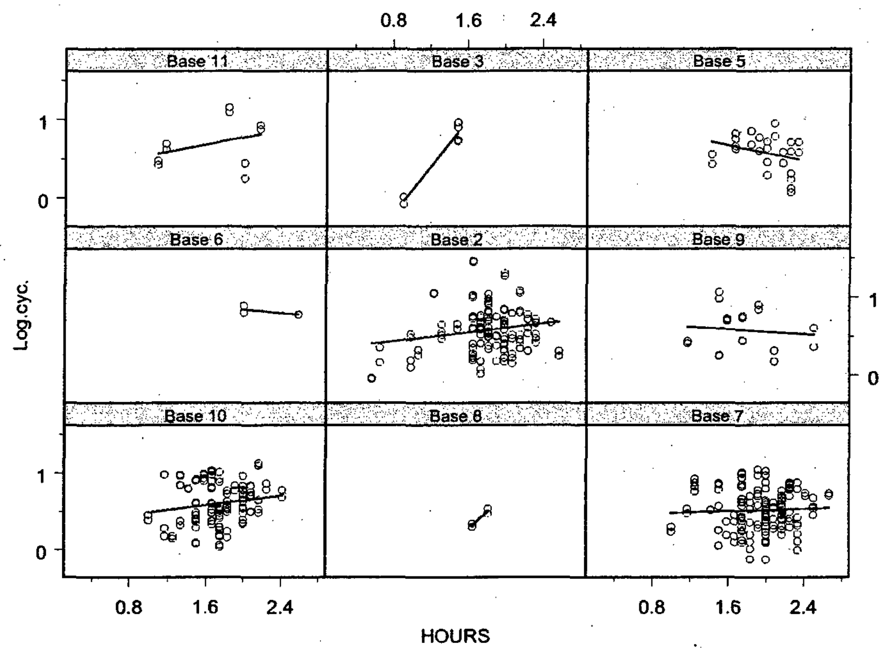

Figure 5 Regression slope of mission duration against log cycles for individual bases

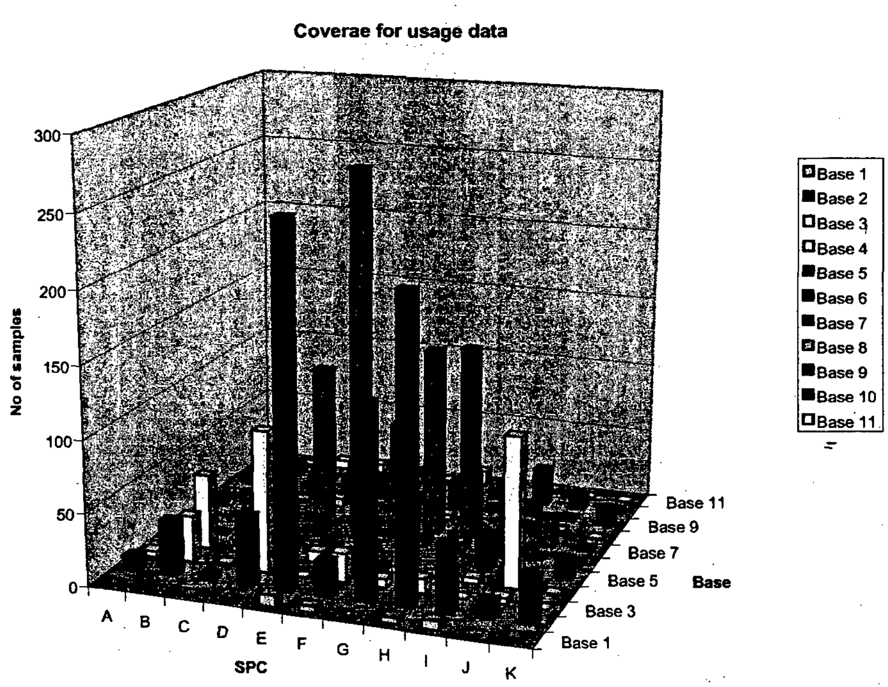

Figure 6 Coverage of data across classification

```
                Df Sum of Sq  Mean Sq  F Value
    LOCATION    10    9.3294 0.932936 14.86182
         SPC    10   23.9976 2.399761 38.22860
       HOURS     1    6.0223 6.022310 95.93643
LOCATION:SPC    64   11.7982 0.184347  2.93668
   Residuals 2258  141.7436 0.062774

                     Pr(F)
    LOCATION 0.000000e+000
         SPC 0.000000e+000
       HOURS 0.000000e+000
LOCATION:SPC 2.131628e-013
   Residuals

Call: lm(formula = Log.cyc. ~ LOCATION + SPC + HOURS + LOCATION * SPC, singular.ok
= T)
Residuals:
     Min     1Q   Median     3Q    Max
 -0.8069 -0.163 0.001212 0.1607 0.9077
```

Figure 7 Output from statistical fitting routine

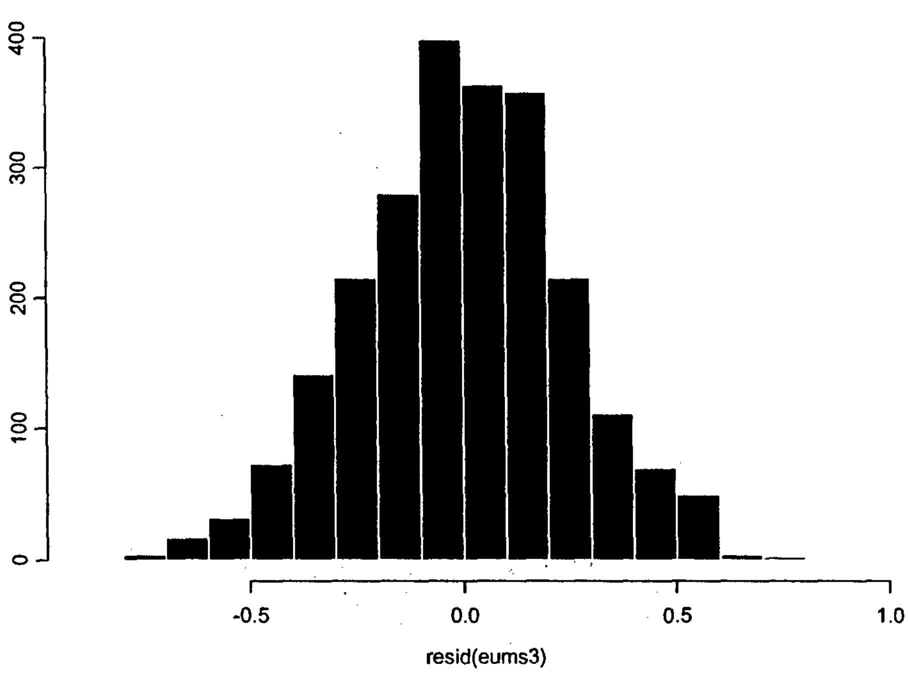

Figure 8 Histogram of residuals

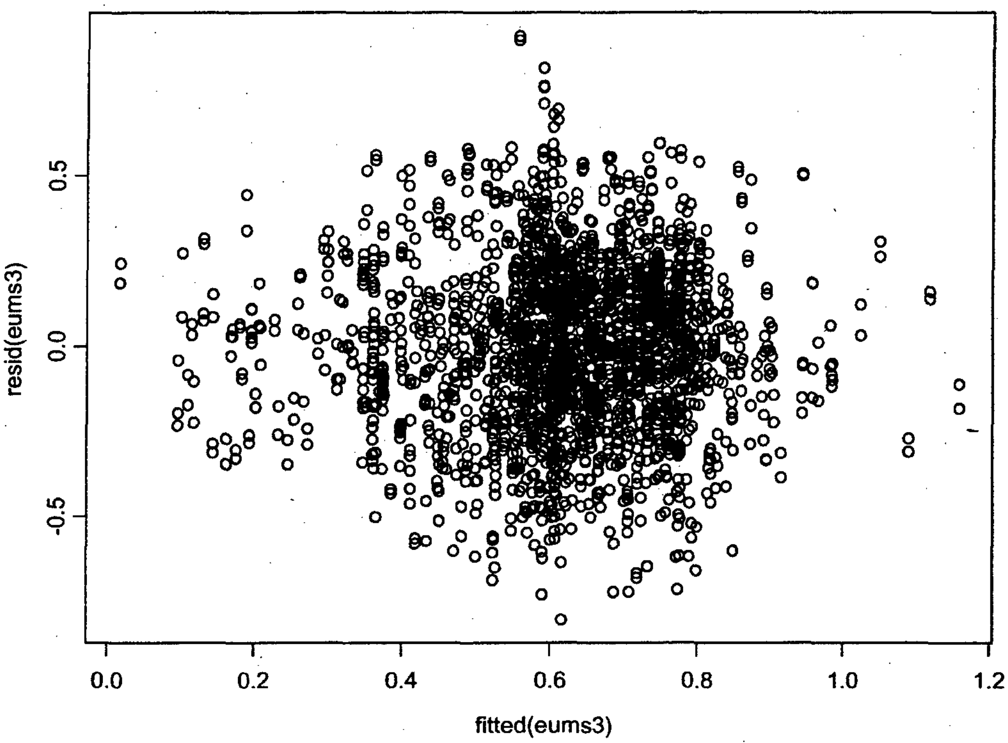

Figure 9 Plot of residuals

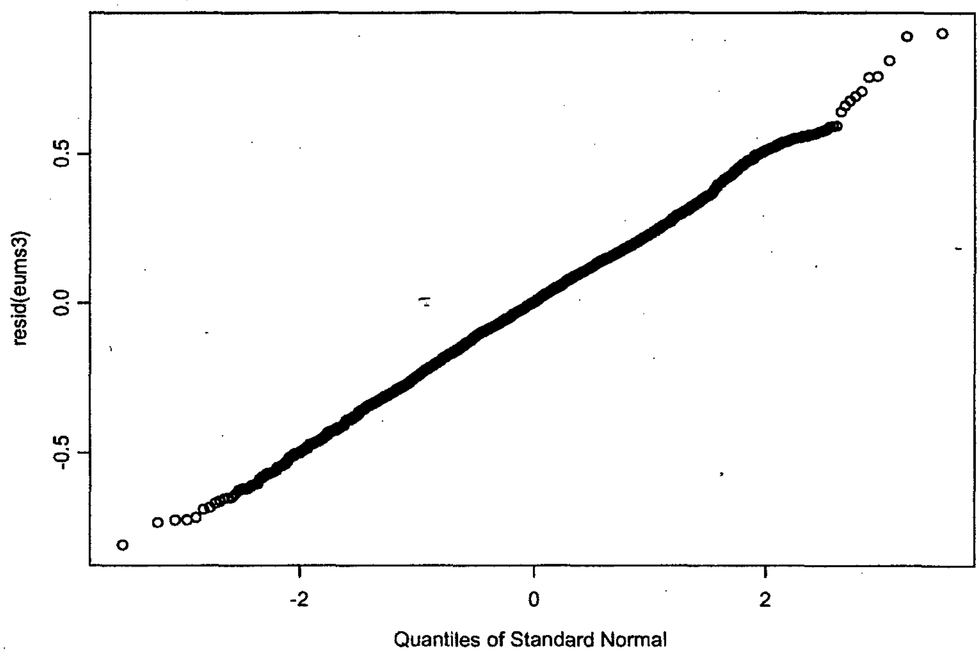

Figure 10 Quantile-quantile plot of residuals

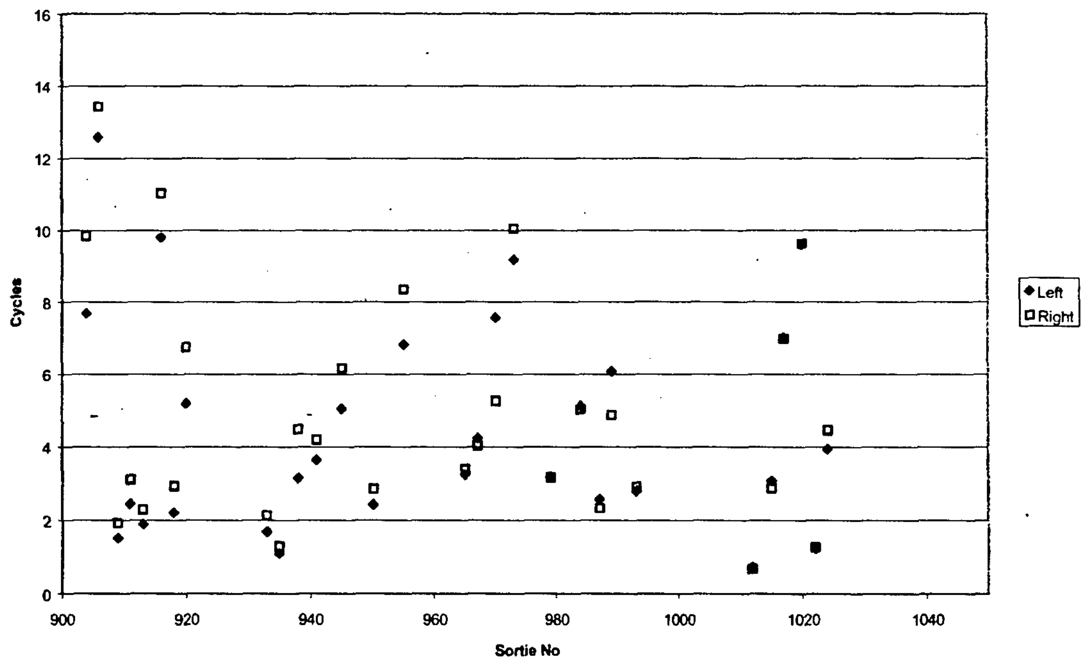

Figure 11 Relative damage of left and right engines

Paper, 'The Assessment of Engine Usage Data': Discussion

Question from S Mosset – SNECMA, France

What is the sortie pattern code and how does the cycle counting operate?

Presenter's Reply

The sortie pattern code is a numerical index that refers to a categorisation of mission types. The cycle counting is done using damage algorithms run post-flight on ground-based computers.

Question from H Pfoertner – MTU, Germany

I would like to make the following observations, based on our experience of GAF operations, in support of your paper:

- The correlation between the duration of a single flight and life usage at a selected critical area is, in most cases, weak. We prefer to use damage per flight instead of damage per flight hour.
- There are known reasons for systematically different life usage between the left and right engines in one aircraft. One reason may be due to differences in engine health, requiring different spool speeds to achieve equal thrust. Another may be caused by different operating procedures between the two engines, such as one engine always being started first, introducing additional thermal stresses.
- Mission type/sortie pattern code may translate into quite different usage. Even if the same pilot flies the same mission, with the same aircraft, on two different days, a factor of two in life usage figures has been observed for certain critical areas.

Presenter's Reply

Agreed.